

Listing of Claims

- 1 1. (Original) A method of reducing the number of times a tree data structure is rebalanced
2 comprising the steps of:
 - 3 (a) allowing a sub-tree of the data tree structure to grow unbalanced to a threshold
4 level greater than one; and
 - 5 (b) rebalancing the data tree structure when the threshold level is reached.
- 1 2. (Original) The method of claim 1 wherein the threshold level is $\log_2 n$ for a tree data
2 structure having about n nodes.
- 1 3. (Original) The method of claim 1 wherein the threshold level is a constant number of levels
2 greater than a level of a balanced portion of the tree data structure.
- 1 4. (Original) The method of claim 1 wherein the step of rebalancing the tree data structure
2 further comprises:
 - 3 (a) developing first and second sets of rebalancing operation tasks, the first set of
4 operation tasks operable to effect a first set of element state transitions and the
5 second set of operation tasks operable to effect a second set of element state
6 transitions, the first and second set of element state transition being distinct one
7 from the other;
 - 8 (b) performing the first set of operation tasks in a first phase; and
 - 9 (c) performing the second set of operation tasks in a second phase.

1 5. (Original) A method of deferring the rebalancing of a tree data structure comprising the
2 steps of:

3 (a) allowing a sub-tree of the tree data structure to grow unbalanced to a length
4 greater than one; and

5 (b) rebalancing the tree data structure when the length of the sub-tree reaches a
6 threshold level.

1 6. (Original) The method of claim 5 wherein the threshold level is $\log_2 n$ for a tree data
2 structure having about n nodes.

1 7. (Original) The method of claim 5 wherein the threshold level is a constant number of levels
2 greater than a level of a balanced portion of the tree data structure.

1 8-9. (Cancelled)

1 10. (Original) A method of performing a rebalancing operation upon a tree data structure
2 comprising the steps of:

3 (a) allowing a sub-tree of the tree data structure to grow unbalanced to a threshold
4 level greater than one;

5 (b) developing, in the case where the sub-tree reaches the threshold level, first and
6 second sets of rebalancing operation tasks, the first and second set of rebalancing
7 operation tasks operable to effect a first and second set of element state transitions
8 respectively;

9 (c) performing the first set of operation tasks in a first phase; and

10 (d) performing the second set of operation tasks in a second phase.

1 11-12. (Cancelled)

1 13. (Original) A system for deferring the rebalancing of a tree data structure comprising:

2 (a) a memory for storing the tree data structure; and

3 (b) a processor coupled to the memory, the processor operable to track the
4 performance of operations upon the tree data structure and rebalance the tree data
5 structure when an unbalanced sub-tree of the tree data structure reaches a
6 threshold level greater than one.

1 14. (Cancelled)

1 15. (Currently Amended) A system for ~~deferring the rebalancing of a tree data structure~~
2 comprising:

3 ~~(a)~~a means for storing the a tree data structure; and

4 a means for tracking the ~~performance~~ execution of operations upon the tree data
5 structure; and

6 ~~and means for~~ rebalancing the tree data structure when an unbalanced sub-tree of the tree
7 data structure reaches a threshold level greater than one, the rebalancing including
8 a first rebalancing phase in which rebalancing operations are executed in parallel
9 and nodes of the unbalanced sub-tree are unlocked, and a second rebalancing
10 phase in which different rebalancing operations are executed.

1 16. (Cancelled)

1 17. (Original) A method of deferring the rebalancing of a tree data structure comprising

2 the steps of:

3 (a) tracking the performance of operations upon the tree data structure; and

4 (b) rebalancing the tree data structure when an unbalanced sub-tree of the tree

5 data structure reaches a threshold level greater than one, the rebalancing

6 further comprising creating a first set of rebalancing operation tasks, the

7 first set of rebalancing operation tasks being characterized by navigation

8 of the tree data structure using at least an existing link, creating a second

9 set of rebalancing operation tasks, the second set of rebalancing operation

10 tasks being different from the first set of rebalancing operation tasks and

11 being characterized by location of elements within the tree data structure

12 using at least one pointer created by the first set of rebalancing operation

13 tasks, and performing at least one operation task of the first set of

14 rebalancing operation tasks in a first phase and at least one of the second

15 set of rebalancing operation tasks in a second phase.

1 18. (Original) A method of deferring the rebalancing of a tree data structure comprising

2 the steps of:

3 (a) tracking the performance of operations upon the tree data structure; and

4 (b) rebalancing the tree data structure when an unbalanced sub-tree of the tree

5 data structure reaches a threshold level greater than one, the rebalancing

6 further comprising executing simultaneous rebalancing operations on the

7 tree data structure including performing any first phase operation task of

8 each of the simultaneous rebalancing operations in a first phase using
9 parallel processes, developing a set of serial rebalancing operations during
10 the first phase, and performing any second phase operation task of each of
11 the simultaneous rebalancing operations in a second phase, the second
12 phase operation task having at least one of the set of serial rebalancing
13 operations.

1 19. (New) A method of rebalancing a tree data structure, the method comprising:
2 allowing a sub-tree of the tree data structure to grow unbalanced until a threshold
3 level is reached;
4 developing a first set of rebalancing operation tasks, the first set of operation tasks
5 operable in parallel on one or more unlocked nodes of the tree data
6 structure during a first phase of the rebalancing;
7 developing a second set of rebalancing operation tasks during execution of the
8 first set of rebalancing operation; and
9 executing the second set of rebalancing operation tasks during a second phase of
10 the rebalancing.

1 20. (New) The method of claim 19, wherein execution of the second set of rebalancing
2 operation tasks is performed without navigating between nodes of the sub-tree.

1 21. (New) The method of claim 19, wherein execution of the first set of rebalancing
2 operations includes generating a list of pointers to nodes requiring updating in the
3 second phase of rebalancing.

1 22. (New) A method of rebalancing a tree data structure, the method comprising:
2 allowing a sub-tree of the data tree structure to grow unbalanced until a threshold
3 level is reached;
4 executing a first set of rebalancing operation tasks during a first rebalancing
5 phase, the first rebalancing phase being characterized by navigation
6 between nodes of the sub-tree; and
7 executing a second set of rebalancing operation tasks during a second rebalancing
8 phase, the second rebalancing phase including navigation to two or more
9 nodes of the sub-tree, the navigation being independent of pointers
10 between nodes of the sub-tree.

1 23. (New) The method of claim 22, wherein the first set of rebalancing operation tasks
2 are performed on unlocked nodes of the sub-tree.

1 24. (New) The method of claim 22, wherein the first set of rebalancing operation tasks
2 includes a plurality of operation tasks configured for parallel execution.

1 25. (New) A method of maintaining a tree data structure, the method comprising:
2 allowing the tree data structure to grow unbalanced;
3 performing first set of rebalancing operation tasks during a first rebalancing phase
4 on a plurality of nodes in the tree data structure, the first set of rebalancing
5 operation tasks being configured for execution while the plurality of nodes
6 are unlocked and for insertion and deletion of nodes; and

7 performing a second set of rebalancing operation tasks on the plurality of nodes in
8 a second rebalancing phase, the second set of rebalancing operation tasks
9 being different than the first set of rebalancing operation tasks and being
10 configured for further operations on the plurality of nodes, the second
11 rebalancing phase occurring after completion of the first rebalancing
12 phase.

1 26. (New) The method of claim 25, wherein the first set of rebalancing operation tasks
2 are performed in parallel.